MoDAVE: Analyzing and visualizing Mosaic climate data

A. Pletzer[1], P. Hamill[1], S. Muszala[1], D. Williams[2], and C. Doutriaux[2]
[1]Tech-X Corporation; [2]Lawrence Livermore National Laboratory

Summary

Climate models are actively moving away from longitude-latitude based grids in order to overcome the difficulties associated with numerical singularities at the north and south poles. Mosaic grids are a new class of grids that combine the advantages of longitude-latitude grids, in terms of numerical efficiency, with the benefits of unstructured grids, in terms of avoiding pole-like singularities. The cubed-sphere grid used by Finite Volume (FV) atmospheric dynamic cores are an example of a Mosaic grid comprising of six structured grids that are assembled to form a nearly regular tiling of the sphere. Presently, there are no general software tools capable of understanding and exploiting the complex connectivity information that characterize cubed-sphere and other Mosaic grids. To remedy this situation, Tech-X and LLNL are developing software called the Mosaic Data Analysis and Visualization Extension (MoDAVE) that will enable other data analysis tools (e.g. CDAT) to read and interpret Mosaic data. CDAT is one of several standard analysis tools that are needed to operate on the Earth System Grid's (ESG's) "Product Service" back-end.

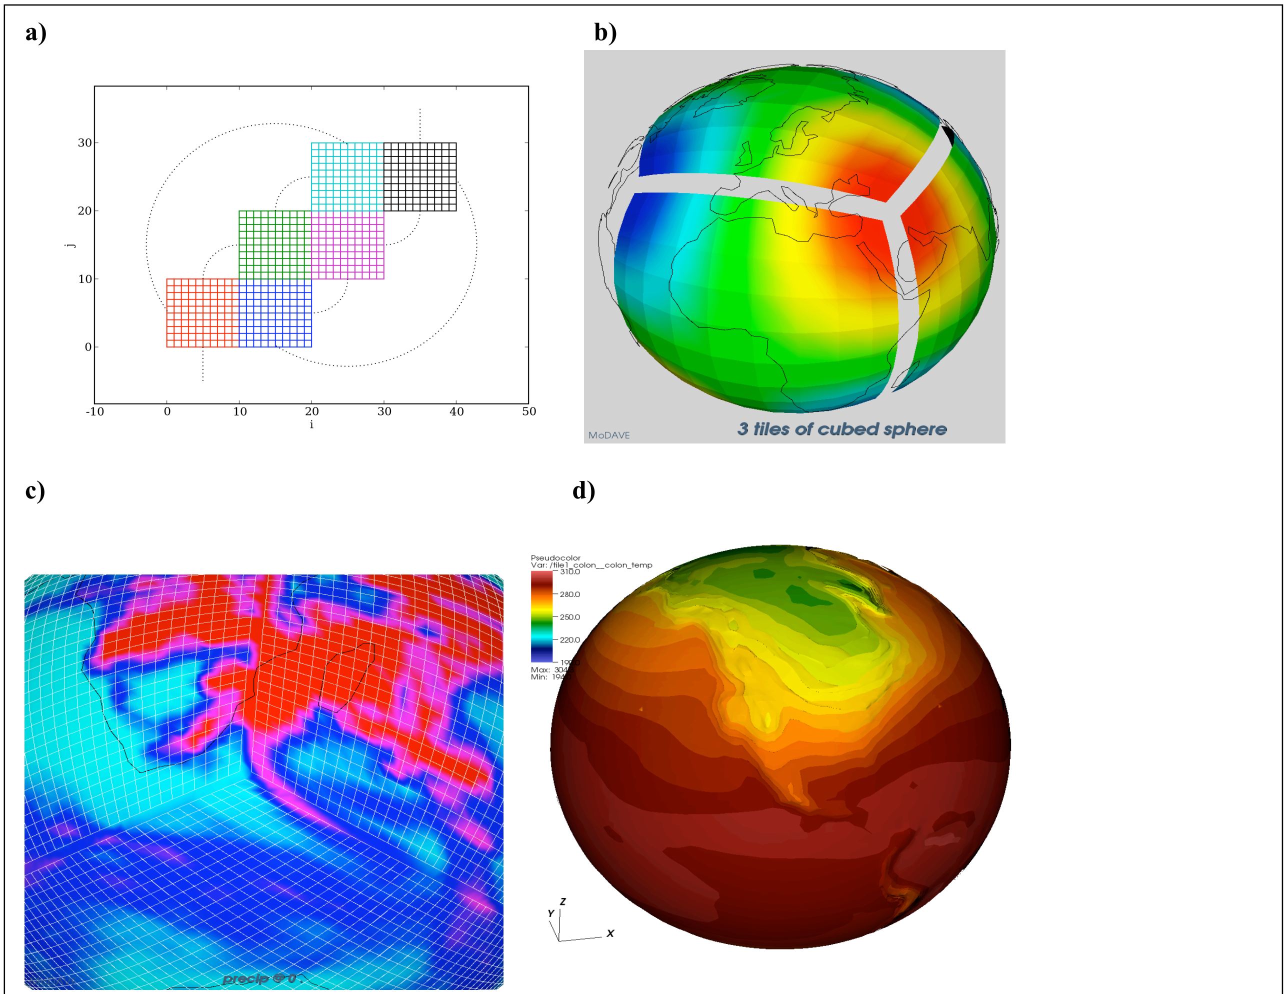

Figure 1. Examples of a Mosaic Grid: Cubed Sphere Connectivity.

With the Intergovernmental Panel on Climate Change 5th Assessment Report (AR5) campaign having started, data providers will no longer be required to supply data on longitude-latitude grids. To prepare the Earth System Grid's (ESG's) Program for Climate Model Diagnosis and Intercomparison (PCMDI) portal to accept simulation data on native grids, we are developing MoDAVE to help third party software infer the Mosaic connectivity patterns of cubed-sphere and other types of grids.

Mosaic grids are becoming increasingly attractive as a way to overcome the severe Courant-Friedrichs-Lewy (CFL) stability condition imposed by longitude-latitude grids, a problem that is exacerbated by the demands on climate models to further improve their resolution. Examples of Mosaic grids include the tri-polar grid implemented by some ocean models, the structured grid recently proposed by Calhoun, Helzel, and LeVeque [SIAM Review 50 (2008), 723-752], and the cubed-sphere grid used in finite volume (FV) atmospheric dynamical cores.

All of these grids share the common characteristic of involving one or more structured grids with a non-standard folding pattern. In the case of the tri-polar grid, folding-type boundary conditions apply along the shortest line connecting the two land poles – the cubed sphere grid in Figure 1a exhibits a more complex folding pattern. Failing to take into account the proper connectivity of nodes across a cut or fold can lead to spurious post-processing or visualization effects. For instance, visualizing cell-centered data on a cubed-sphere mesh without using the connectivity information yields gaps between tiles, as shown in Figure 1b.

MoDAVE, funded by a Phase I Small Business Innovation Research (SBIR), has explored different approaches for accurately rendering Mosaic datasets. In Figure 1c, we show how the connectivity information can be exploited to fill in the gaps of Figure 1b. The figure shows precipitation data southeast of the southern African continent tip, a region where three tiles of the cubed-sphere merge. By creating a structured grid that is one-cell wide and spans the length of a tile, the data on neighboring tiles can be bridged, giving the impression of smooth data transition across tiles. However, this would still leave a small opening in the form of a triangular cell at each of the eight points on the sphere where three tiles join. Filling theses eight corner points with a triangular cell now produces the desired effect.

MoDAVE is a visualization agnostic tool – independent of the visualization framework. The code has been written with the aim of allowing users to interchange visualization engines as they see fit. Presently, we support a raw VTK engine, which is invoked through a Python interface and was used to produce Figure 1b and 1c. Alternatively, users can choose to save their data in an HDF5 file, which complies with the VizSchema format recently developed for scientific visualization at Tech-X. A ViSit plug-in exists that reads in VizSchema and we used this approach to produce Figure 1d.

Another area of active research involves exploiting the natural parallelism of Mosaic datasets. Typically, the data on each tile are stored in different files. In order to allow MoDAVE to scale with increased model resolutions we experimented with several approaches for accelerating post-processing and visualization tasks. One approach uses task farming. Another approach relies on the development of a distributed array class in Python, which uses the MPI-2 standard for communication. A third approach involves moving computationally intensive operations of MoDAVE to a Graphical Processing Unit (GPU). GPUs can improve computational performance by an order of magnitude or more. We believe that these three approaches will enable post-processing tools using MoDAVE (e.g., CDAT) to remain responsive in the face of future challenges, the most important of which is how to handle the massive increase in data that will occur as models improve their resolution.

For further information on this subject contact:

- Name: Alexander Pletzer
  Organization: Tech-X
  Email: pletzer@txcorp.com
- Name: Dean N. Williams.
  Organization: PCMDI, LLNL
  Email: williams13@llnl.gov